

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

The Varieties of Dynamicism

Permalink

<https://escholarship.org/uc/item/9j96c8hr>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 31(31)

ISSN

1069-7977

Author

Zednik, Carlos

Publication Date

2009

Peer reviewed

The Varieties of Dynamicism

Carlos Zednik (zednik@indiana.edu)
Cognitive Science Program, Indiana University
Bloomington, IN 47406 USA

Abstract

The dynamical approach to cognition is often considered to be ‘revolutionary’. In contrast to the well-established approaches of computationalism and connectionism, dynamicism is typically thought to be anti-representational, holistic, phenomenological and law-based. In this paper, I will argue that this way of thinking about dynamicism is too restrictive: it fails to capture the heterogeneous nature of dynamicist research. Although all dynamicist research projects share a commitment to the mathematical methods, tools, and concepts of dynamical systems theory, they frequently disagree with respect to the truth or falsity of representationalism, the role of holistic phenomenological modeling, and in general, the nature of dynamical explanation.

Keywords: dynamical systems theory; dynamicism; explanation; representation.

The Standard Conception of Dynamicism

The first comprehensive philosophical treatment of the dynamical approach to cognition is due to Tim van Gelder (van Gelder 1995; 1998; 1999; van Gelder & Port 1995). van Gelder aimed to characterize dynamicism as a unified body of research, and to distinguish it from the established research programs of computationalism and connectionism. To this end, he provided a comprehensive overview of the methods, tools, and concepts of dynamical systems theory—the *mathematical* framework at the heart of the approach. In addition, he promoted a particular set of explanatory principles which, he suggested, constitute the *theoretical* framework that unifies all branches of dynamicist research. Although sometimes left implicit by van Gelder himself, this particular theoretical framework—what I call the *standard conception of dynamicism*—has by now developed into a consensus view about the theoretical commitments of dynamicist researchers.

Before discussing the standard conception in detail, it is worth briefly considering the principal motivations for the dynamical approach. Increasingly, it is becoming clear that cognitive phenomena arise from reciprocal nonlinear interactions between brain, body, and environment (Beer 1995; Clark 1997). Such interactions, it has been argued (Chemero & Silberstein 2008; van Gelder 1995; Wheeler 2005), can be difficult to accommodate within computationalism and connectionism. Although for the purposes of this paper it will not be necessary to determine the fate of these traditional approaches, it will be important to recognize that dynamicism is frequently considered to be a uniquely viable alternative. Why?

According to the standard conception, dynamicism promises to adequately describe reciprocal nonlinear interactions between brain, body, and environment by

adhering to the explanatory principles of *anti-representationalism* (Chemero 2000; Thelen & Smith 1994; van Gelder 1995) and *holism* (Chemero & Silberstein 2008; Wheeler 2005). Although it has already been shown that the mere presence of reciprocal nonlinear interactions does not logically entail the absence of representations (Bechtel 1998; Clark 1998), proponents of the standard conception typically find “the notion of representation to be dispensable or even a hindrance for their particular purposes” (van Gelder 1998:622). To a large extent, this is because the presence of such interactions seems to preclude the possibility of identifying structurally and functionally isolated ‘representation-producers’ and ‘representation-consumers’ (Wheeler 2005). Instead, dynamicists describe a cognitive system’s behavior holistically, in terms of differential equations that are defined over a small number of collective variables and global control parameters.

What exactly do such equations describe, and what is the nature of the explanations they provide? Because the physical structure of nonlinear brain-body-environment systems can be heterogeneous and complex, collective variables and global control parameters are frequently used to capture the observable structure of behavior itself, rather than the physical structure of the system that underlies such behavior. That is, dynamical models resemble what physicists sometimes call *phenomenological models*—models that directly relate measurable quantities to one another (Beek et al. 1995; Cartwright 1983). Notably, when a particular model captures a wide variety of related phenomena, and can be used to derive predictions about them in factual and counterfactual circumstances, it is said to instantiate a general principle or law that explains those phenomena. That is, according to the standard conception of dynamicism, dynamical explanations of cognitive phenomena take the form of *covering law explanations* (Bechtel 1998; van Gelder 1999; Walmsley 2008).

Thus construed, dynamical models can be contrasted with computational and connectionist *process models*, which aspire to be more-or-less abstract descriptions of the physical and functional processes from which observable behaviors arise (Marr 1982; Luce 1995). Process modeling involves making explicit certain assumptions about structural and functional organization, including assumptions about possible ways to functionally *decompose* the target behavior, and about how individual functional components might be *localized* in the system from which the target behavior arises (Bechtel & Richardson 1993; Cummins 1982; 2000). In virtue of relying on process models rather than phenomenological models, computational and connectionist explanations differ from

dynamical explanations as the standard conception understands them: they resemble *mechanistic explanations* that involve the description of organized collections parts and operations that give rise to the phenomena being explained (Bechtel 1998; 2007; Machamer, Darden & Craver 2000).

In this paper, I will not question the explanatory power of the theoretical framework that embraces the explanatory principles of anti-representationalism, holistic phenomenological modeling, and covering law explanation. Indeed, I acknowledge that this framework is likely to be exceedingly powerful when it comes to explaining the behavior of complex, nonlinear, cognitive systems that span the boundaries between brain, body, and environment. Nevertheless, at least as important as the question of explanatory power is the question of empirical relevance: to what extent does the theoretical framework offered by the standard conception actually play a role in contemporary dynamicist research?

To answer this question, I will distinguish three families of dynamicist research. Of those, only one—*Kelsonian dynamicism*—adheres to the explanatory principles of anti-representationalism, holistic phenomenological modeling, and covering law explanation. In contrast, the other two families of dynamicist research—*dynamical field theory* and *dynamical agent modeling*—appear to conflict with, or at least remain agnostic about, the theoretical framework offered by the standard conception. For this reason, I suggest that the standard conception of dynamicism underestimates the heterogeneous nature of the dynamical approach to cognition. Rather than collectively embracing a particular set of explanatory principles, the families of research that together make up the dynamical approach are unified only by their shared commitment to the methods, tools, and concepts of dynamical systems theory.

The Varieties of Dynamicism

Kelsonian Dynamicism

I will begin by discussing a family of dynamicist research that generally adheres to the theoretical framework offered by the standard conception, and that (possibly as a result) is often assumed to be representative of dynamicist research in general. The *Kelsonian* family of research is modeled after J.A. Scott Kelso's work on coordination dynamics (Kelso 1995). Characteristically, Kelsonian research projects explain perception/action phenomena by specifying a small set of differential equations defined over an equally small set of collective variables and global control parameters. For example, the well-known Haken-Kelso-Bunz (HKB) model of bimanual coordination captures the motion of two rhythmically oscillating index fingers with the following dynamical model:

$$\dot{\phi} = -a \sin \phi - 2b \sin 2\phi$$

In this model, the changing state of the between-fingers phase relation, ϕ , is expressed as a function of the oscillation frequencies of two moving fingers, b/a . The

model accurately predicts a qualitative change in the behavior of ϕ whenever b/a moves from low values to high. Whereas at low frequencies ϕ settles on either one of two phase relations (in-phase or anti-phase), at high frequencies it reliably settles on the in-phase relation only. In dynamical terms, the model describes a one-dimensional dynamical system with two point-attractors when the value of parameter b/a is low, and undergoes a bifurcation in which the two point-attractors coalesce into one when that parameter reaches a certain critical value.

Kelsonian dynamicism reinforces the standard conception. Although the phenomenon of bimanual coordination arises from the (presumably) nonlinear interactions between motor neurons, muscle tissue, and environmentally-situated index fingers, ϕ and b/a each correspond to measurable features of the phenomenon itself. Moreover, the differential equation at the heart of the HKB model is not derived from assumptions about the physical and functional organization of the system from which bimanual coordination emerges, but is merely the simplest equation that adequately accounts for the observed data (Kelso 1995). In this sense, the HKB model is a paradigmatic phenomenological model. Interestingly, although Kelso repeatedly asserts that the HKB model avoids the need to hypothesize a neural 'switching mechanism', it is also perfectly compatible with such a hypothesis—the HKB model is a phenomenological model that is consistent with any number of process accounts of the structure of the underlying system.

The HKB model is also holistic and non-representational. The equation at the heart of the model describes the behavior of the bimanual coordination system as a whole, without explicitly identifying its individual components: ϕ is a collective variable that captures a *relation* between two moving fingers, and b/a is a global control parameter that determines the attractor landscape of the system as a whole. Moreover, because ϕ and b/a each correspond to behavioral properties of bimanual coordination rather than to structural and functional properties of the system from which that behavior arises, neither one describes "internal information-carrying states of an organism" (Dietrich & Markman 2001: 332). Although internal information-carrying (i.e. potentially representational) states may in fact be involved in the mechanism that underlies bimanual coordination, the presence of such a mechanism would have to be independently posited and confirmed. In Chemero's (2000) words, the HKB model allows us to take a thoroughly anti-representational "dynamical stance" toward the phenomenon of bimanual coordination.

The HKB model is a holistic, non-representational, and phenomenological model of bimanual coordination. But how does it *explain* bimanual coordination? Given a particular set of initial conditions (the initial phase relation and oscillation frequency), the equation at the heart of the HKB model can be used to derive predictions about the future state of the system (e.g. to determine the particular phase relation at which it will eventually settle, or to

calculate how long individual phase-transitions will take). Moreover, although the HKB model was originally developed to explain the phenomenon of bimanual coordination, variations of the same model have been used to account for rhythmic stress-patterns in regular speech (Port 2003) and ‘informational’ coupling of rhythmic limb motion between individuals (Ouille et al. 2005). In other words, since the derivation of predictions is the primary purpose of the HKB model, and since the regularity captured by that model appears to be a general principle that applies to a variety of related phenomena, it is in line with the explanatory commitments of the standard conception: Kelsonian dynamicism strives to provide covering law explanations (see also Walmsley 2008).

This brief discussion of the HKB model suggests that the Kelsonian family of dynamicist research does in fact adhere to the explanatory principles promoted by the standard conception of dynamicism. Notably, this kind of research is often considered paradigmatic (see e.g. Chemero & Silberstein 2008). Nevertheless, in what follows, I will suggest that the Kelsonian family of research is not in fact all that representative of dynamicist research in general. In particular, I will argue that unlike Kelsonian dynamicism, at least two prominent families of dynamicist research largely reject or remain relatively agnostic about explanatory principles of anti-representationalism, holistic phenomenological modeling, and covering law explanation.

Dynamical Field Theory

A second prominent family of dynamicist research is *dynamical field theory* (DFT). A well-known member of this family is Thelen et al.’s (2001) model of infant perseverative reaching in the A-not-B task:

$$\tau \dot{u}(x,t) = -u(x,t) + S(x,t) + g[u(x');x']$$

This model specifies the activation values of a high-dimensional ‘motor planning field’ (u) that depends on the field’s previous activation ($-u$); a series of inputs (S) that correspond to the changing and unchanging features of the task environment as well as a memory trace of previous reaches; the level of cooperative interaction (g) between individual points on the field; and a temporal decay constant (τ). Every point in the motor planning field corresponds to a particular spatial location in the A-not-B task environment. If at any moment the activation value at a single point increases beyond a certain threshold value, a reach is induced toward the corresponding location. Crucially, the likelihood that the field’s activation value surpasses the threshold is a function of the ‘cooperativity parameter’ g : at high g -values, the activation of every point in the field is positively influenced by the activation of its neighbors, thus increasing the probability that accurate reaches are induced, and allowing the field to maintain stable activation levels for a period of time even in the absence of immediate sensory input. Psychologically, g corresponds to the parameter that determines whether or not accurate goal-directed reaches can be performed at different stages of an infant’s development.

For current purposes, the most significant fact about Thelen et al.’s DFT model is that the value of the input vector S is a function of three independent vectors: a ‘task input’ which captures the unchanging features of the A-not-B task environment; a ‘specific input’ which corresponds to the changing perceptual scene in each trial; and a ‘memory trace’, which corresponds to the effect of remembered reaches from earlier trials. It is significant for at least two reasons.

First, Thelen et al.’s motivation for defining S in this particular way is independent of infants’ actual performance in the A-not-B task—nothing in the *data* indicates that the equation at the heart of their model should reflect the distinction between specific, task, and memory inputs. Rather, Thelen et al.’s primary motivation for defining S as they do is an assumption about the functional structure of the infant movement planning system—the assumption that the system from which infant perseverative reaching involves three functionally separable components. In other words, unlike the HKB model, Thelen et al.’s DFT model is a process model that incorporates a variety of assumptions about the structure of the system from a particular cognitive phenomenon arises.

Second, Thelen et al.’s definition of S in terms of three separable input sources is in effect a *functional analysis* of goal-directed reaching. Functional analysis (Cummins 1983; 2000) involves decomposing a complex function, capacity, or behavior P into a set of simpler functions, capacities or behaviors $p_1 \dots p_n$ that work together to produce P . The analysis of movement planning into the relative contributions of task input, specific input, and memory trace is an analysis of exactly this kind. Therefore, although the model’s ‘cooperativity parameter’ g is a global control parameter that constrains the behavior of the system as a whole, there is a clear sense in which Thelen et al. are also invoking the explanatory principle of decomposition—a principle more commonly associated with computationalism and connectionism than with dynamicism (Bechtel 1998; Cummins 2000).

The fact that Thelen et al.’s DFT model conflicts with the standard conception’s commitment to holistic phenomenological modeling suggests that it might also deviate from the commitment to covering law explanation. Walmsley (2008) explicitly defends the view that Thelen et al.’s DFT model provides a covering law explanation of infant perseverative reaching, but this view appears to be mistaken. Although the equation at the heart of the model can be used to derive predictions about the future state of the system (e.g. the direction of reach) from a particular set of initial conditions (e.g. initial field activation values and a particular set of inputs), and although the model has been adapted to account for a wide variety of movement-planning phenomena, interpreting it as a widely-applicable predictive instrument is to misunderstand the intentions of the authors. Rather than seeking to uncover general principles or laws of movement planning *in general*, Thelen et al. are in the business of describing the specific processes that underlie

movement planning *in humans*. Accordingly, they even go so far as to appeal to a preliminary neural localization:

“At this point, we conceptualize this field only in abstract terms as a site where visual input and memory are integrated into movement parameters supporting movement amplitude, direction, or time. Later in the discussion, we will speculate further as to possible neuroanatomical areas where such a field might evolve.” (Thelen et al. 2001: 16)

In summary, Thelen et al. are committed not only to describing (via functional analysis) the process that underlies episodes of goal-directed reaching, but are additionally committed to eventually localizing that process in the human brain. Thus construed, Thelen et al. are quite clearly seeking a mechanistic explanation, rather than a covering law explanation, of infant perseverative reaching.

Finally, what about representation? Spencer & Schöner (2003), two of the most prominent contributors to dynamical field theory research, have explicitly construed that family of research as a way of “bridging the representational gap” in the dynamical approach to cognition. Their suggestion is to think of dynamical fields as describing the large-scale neural activation patterns that represent the continuous spatial dimensions of our immediate environment, even when that environment is temporarily occluded or otherwise inaccessible. Of course, it remains to be seen whether or not such a representational construal “earns its explanatory keep” (Ramsey 1997)—i.e. whether it identifies a substantive and explanatorily useful notion of representation. Nevertheless, the fact that this construal is made explicit by two of the most prominent dynamical field theorists suggests that this particular family of research is not after all committed to the anti-representationalism of the standard conception. This observation, combined with the observation that Thelen et al.’s DFT model is also not committed to the explanatory principles of holistic phenomenological modeling and covering law explanation, suggests that the standard conception of dynamicism misrepresents at least this particular family of dynamicist research.

Dynamical Agent Modeling

A third well-established family of dynamicist research—dynamical agent modeling—involves applying the methods, tools and concepts of dynamical systems theory to study the behavior of simulated and artificially evolved brain-body-environment systems (Beer 1995; 1996; 2003; Harvey et al. 2005). A particularly prominent example of this kind of research is Randy Beer’s (2003) model of ‘visual’ categorization in an artificially evolved brain-body-environment system:

$$\begin{aligned} \tau_i \dot{s}_i &= -s_i + I_i(x, y; \alpha) \quad i = 1, \dots, 7 \\ \tau_i \dot{s}_i &= -s_i + \sum_{j=1}^7 w_{ji} \sigma(g(s_j + \theta)) + \sum_{j=8}^{12} w_{ji} \sigma(s_j + \theta_j) \quad i = 8, \dots, 12 \\ \tau_i \dot{s}_i &= -s_i + \sum_{j=8}^{12} w_{ji} \sigma(s_j + \theta) \quad i = 13, 14 \\ \dot{x} &= 5(\sigma(s_{13} + \theta_{13}) - \sigma(s_{14} + \theta_{14})) \\ \dot{y} &= -3 \end{aligned}$$

This 16-dimensional dynamical model describes a simulated agent—equipped with a 14-neuron continuous-time neural network ‘brain’ (neural parameters w , τ , σ , θ) and seven ‘visual’ sensors (input vector I)—that was evolved to categorize objects according to their shape. Two kinds of objects, circles and diamonds, fall towards the agent (at rate y), which must make a categorical discrimination by moving horizontally (at rate x) to catch circles and avoid diamonds. In order to accomplish this task, the agent uses a particularly interesting ‘active scanning’ strategy: it repeatedly moves from side to side to ‘foveate’ the falling object before eventually settling on a position at which it will either catch or avoid. Since this ‘active scanning’ behavior emerged unexpectedly from the artificial evolutionary process, it constitutes an interesting and non-trivial target for dynamical explanation.

How does Beer go about explaining categorization via ‘active scanning’ in this particular simulated brain-body-environment system? Unlike the previously discussed dynamical explanations of bimanual explanation and infant perseverative reaching, here the explanatory burden is not carried by the dynamical model alone. Rather, Beer’s dynamical explanation centers on an extensive *dynamical analysis* of the simulated brain-body-environment system. Crucially, this dynamical analysis relies on an explicitly stated decompositional strategy:

“...we will decompose the agent–environment dynamics into: (1) the effect that the relative positions of the object and the agent have on the agent’s motion; (2) the effect that the agent’s motion has on the relative positions of the object and the agent.” (Beer 2003: 228)

This decompositional strategy constitutes a significant departure from the standard conception’s commitment to holistic modeling. Although Beer acknowledges that categorization via ‘active scanning’ is a property of the whole simulated brain-body-environment system, explaining how this phenomenon arises requires an understanding of the way in which individual parts of the system—the agent on the one hand and the environment on the other—interact. Notably, this explanatory task is accomplished in a uniquely dynamicist fashion. First, Beer characterizes the effects of every possible ‘visual’ input as a parametric change to the attractor landscape of the agent’s two-dimensional (s_{13} , s_{14}) motor neuron state space. Second, he characterizes the continuously changing state of

the environment (defined as the relative positioning of agent and object) as a continuous trajectory through the two-dimensional (x, y) state space. Finally, Beer shows how the observed ‘active scanning’ behavior arises from the interactions between agent and environment by superimposing the continuous (x, y) trajectory on the attractor landscape of the $(s13, s14)$ plane. In general, Beer’s method of decomposing a closed (i.e. autonomous) brain-body-environment system into two mutually interdependent (i.e. non-autonomous) subsystems, as well as his method of describing their reciprocal nonlinear interactions in terms of one subsystem’s effects on the asymptotic behavior of the other, is a particularly vivid demonstration of how the analytic tools of dynamical systems theory can be applied to study not just systems as a whole, but in terms of the reciprocal interactions between multiple simpler components.

Aside from its decompositional nature, Beer’s approach can be contrasted with the theoretical framework of the standard conception in virtue of being process-oriented and indeed, mechanistic. First, a purely phenomenological approach would account for the observed ‘active scanning’ behavior solely in terms of directly observable features of that behavior itself, by identifying its crucial dimension of change (i.e. the agent’s horizontal position), and relating it to the directly observable factors on which that change depends (i.e. the proximity and shape of the falling object). In contrast, the decompositional dynamical analysis outlined above explains the observed behavior by showing how it arises from the changing attractor landscape of the motor neuron state space—the functional structure of the underlying system’s ‘physical’ properties. Second, although Beer’s dynamical model can (in principle) be used to predict the future state of the system from any set of initial conditions, it shouldn’t be thought of as identifying one or more principles or laws of perceptual categorization in general. Indeed, Beer is very explicit about the fact that his intention “is not to propose a serious model of categorical perception”, but rather to examine “in considerable depth each aspect of [the agent’s] behavior, and the mechanisms underlying that behavior” (Beer 2003: 210). In short, Beer’s dynamical explanation is of the mechanistic, rather than the covering law, variety.

What remains to be discussed is the extent of Beer’s commitment to anti-representationalism. The standard conception holds that dynamical agent models cannot or should not be understood in representational terms (Wheeler 2005). Beer himself is somewhat agnostic on this issue, explicitly adopting a stance of “representational skepticism” (Beer 2003). Although he does not in fact invoke representational principles in order to explain the observed behavior, the fact that his analytic strategy involves decomposing a larger system into multiple interacting components makes it amenable to what Chemero & Silberstein (2008) call “representation-hunting”: the practice of identifying distinguishable producers and consumers of information. Contrary to the standard conception of

dynamism, which assumes an *a priori* rejection of representationalism, here the truth or falsity of representationalism is determined *a posteriori*, according to the utility of representational principles for understanding and explaining the behavior of individual dynamical agents.

Conclusion: Methodology or Theory?

The preceding discussion suggests that the standard conception of dynamicism—the conception first outlined by van Gelder and later adopted by most theoretical treatments of the dynamical approach, positive or negative—misrepresents the theoretical commitments of at least two prominent families of dynamicist research. That is, the standard conception underestimates the heterogeneous nature of the dynamical approach to cognition. Although Kelsonian dynamicism does seem committed to the explanatory principles of anti-representationalism, holistic phenomenological modeling, and covering law explanation, dynamical field theory and dynamical agent modeling do not. For this reason, van Gelder’s original goal—to develop a conception of dynamicism as a unified body of research that distinguishes it from both computationalism and connectionism—remains unsatisfied.

In closing, I briefly outline two alternative ways of carving up the logical space of cognitive science research programs that might or might not be more successful. On the first, we retain the original tripartite distinction between computationalism, connectionism, and dynamicism, but accept that the differences are *methodological* at heart: they concern the particular mathematical methods, tools and concepts that practicing cognitive scientists bring to bear to the study of cognition. In particular, each of the three families of dynamicist research discussed above relies on the methods, tools, and concepts of dynamical systems theory. First, they each rely on the practice of *dynamical modeling*: they describe cognitive systems and cognitive phenomena in terms of coupled difference or differential equations. Second, by characterizing state-space trajectories and attractor landscapes, as well as by identifying and classifying critical points and bifurcations, they each make use of the method of *dynamical analysis*. Third, by emphasizing properties like stability, sudden loss of stability, asymptotic behavior and coupled interaction, each one of Kelsonian dynamicism, dynamical field theory, and dynamical agent modeling allows us to study cognitive phenomena from a uniquely *dynamical perspective* that is unlikely to be shared with computationalist and connectionist research projects.

On the second way of individuating cognitive scientific research programs, we retain the ability to distinguish research programs according to their major theoretical commitments, but acknowledge that the traditional distinction between computationalism, connectionism, and dynamicism cuts across important theoretical distinctions. For example, whereas the anti-representational, anti-mechanistic commitments of Kelsonian dynamicism have much in common with classical behaviorism, Gibsonian

ecological psychology, and several strands of contemporary mathematical psychology, the potentially representational, mechanism-oriented research of dynamical field theorists and dynamical agent modelers has more in common with many—though not all—of the explanatory principles of classical cognitivism. On this view, the theoretical foundations of ‘mainstream’ cognitive science need not be upturned in order to accommodate the ubiquity of reciprocal nonlinear interactions between brain, body, and environment—although it may be necessary and instructive to apply the methods, tools, and concepts of dynamical systems theory to articulate a richer and more powerful notion of mechanistic explanation.

References

- Bechtel, W. (1998). Representations and cognitive explanations: Assessing the dynamicist challenge in cognitive science. *Cognitive Science* 22(3):295-317.
- Bechtel, W. (2007). *Mental Mechanisms*. Lawrence Erlbaum.
- Bechtel, W. & Richardson, R.C. (1993). *Discovering complexity: Decomposition and localization as strategies in scientific research*. Princeton: Princeton University Press.
- Beek, P., Peper, C. & Stegeman, D. (1995). Dynamical models of movement coordination. *Human Movement Science* 14:573-608.
- Beer, R.D. (1995). Computational and dynamical languages for autonomous agents. In R. Port and T. van Gelder (Eds.), *Mind as Motion: Explorations in the Dynamics of Cognition* (pp. 121-147). MIT Press.
- Beer, R.D. (1996). Toward the evolution of dynamical neural networks for minimally cognitive behavior. In P. Maes, M. Mataric, J. Meyer, J. Pollack and S. Wilson (Eds.), *From animals to animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior* (pp. 421-429). MIT Press.
- Beer, R.D. (2003). The dynamics of active categorical perception in an evolved model agent (Target article & Reply to comments). *Adaptive Behavior* 11(4):209-243.
- Cartwright, N. (1983). *How the Laws of Physics Lie*. Oxford University Press.
- Chemero, A. (2000). Anti-representationalism and the Dynamical Stance. *Philosophy of Science* 67(4):625-647.
- Chemero, A. & Silberstein, M. (2008). After the Philosophy of Mind: Replacing Scholasticism with Science. *Philosophy of Science* 75:1-27.
- Clark, A. (1997). *Being There: Putting Brain, Body, and World back Together Again*. Cambridge, MA: MIT Press.
- Clark, A. (1998). Time and Mind. *Journal Of Philosophy* XCV(7):354-376.
- Cummins, R. (1983). *The Nature of Psychological Explanation*. Cambridge, MA: MIT Press.
- Cummins, R. (2000). ‘How Does it Work?’ versus ‘What are the Laws?’: Two Conceptions of Psychological Explanation. In F. Keil & R. Wilson (Eds.), *Explanation and Cognition* (pp. 117-144). Cambridge, MA: MIT Press.
- Craver, C. (2007). *Explaining the Brain*. Oxford University Press.
- Dietrich, E. & Markman, A. (2001). Dynamical description versus dynamical modeling: Reply to Chemero. *Trends in Cognitive Sciences* 5(8):332.
- Harvey, I., Di Paolo, E., Wood, R., Quinn, M. & Tuci, E. (2005). Evolutionary Robotics: A New Scientific Tool for Studying Cognition. *Artificial Life* 11:79-98.
- Kelso, J.A.S. (1995). *Dynamic Patterns: The Self-Organization of Brain and Behavior*. Cambridge, MA: MIT Press.
- Luce, D. (1995). Four Tensions Concerning Mathematical Modeling in Psychology. *Annual Review of Psychology* 46:1-26.
- Machamer, P., Darden, L. & Craver, C.F. (2000). Thinking About Mechanisms. *Philosophy of Science* 67:1-25.
- Marr, D. (1982). *Vision*. New York: Henry Holt & Co.
- Oullier, O., de Guzman, G.C., Jantzen, K.J., Lagarde, J.F., & Kelso, J.A.S. (2005). Spontaneous interpersonal synchronization. In C. Peham, W.I. Schöllhorn, & W. Verwey (Eds.). *European workshop on movement sciences: Mechanics-Physiology-Psychology* (pp. 34-35).
- Port, R. (2003). Meter and speech. *Journal of Phonetics* 31:599-611.
- Spencer, J.P. & Schöner, G. (2003). Bridging the representational gap in the dynamic systems approach to development. *Developmental Science* 6:392-412.
- Thelen, E. & Smith, L.B. (1994). *A Dynamic Systems Approach to the Development of Cognition and Action*. Cambridge, MA: MIT Press.
- Thelen, E., Schöner, G., Scheier, C. & Smith, L.B. (2001). The dynamics of embodiment: A field theory of infant preservative reaching. *Behavioral and Brain Sciences* 24: 1-34.
- van Gelder, T.J. (1995). What might cognition be, if not computation? *Journal of Philosophy* 91(7): 345-381.
- van Gelder, T.J. (1998). The dynamical hypothesis in cognitive science. *Behavioral and Brain Sciences* 21: 1-14.
- van Gelder, T.J. (1999). Defending the dynamical hypothesis. In W. Tschacher & J.P. Dauwalder (Eds.), *Dynamics, Synergetics, Autonomous Agents: Nonlinear Systems Approaches to Cognitive Psychology and Cognitive Science* (pp. 13-28). Singapore: World Scientific.
- van Gelder, T.J. & Port, R. (1995). It’s about time: An Overview of the Dynamical Approach to Cognition. In R. Port & T.J. van Gelder (Eds.), *Mind as Motion: Explorations in the Dynamics of Cognition* (pp. 1-44). Cambridge MA: MIT Press.
- Walmsley, J. (2008). Explanation in Dynamical Cognitive Science. *Minds and Machines* 18(3):331-348.
- Wheeler, M. (2005). *Reconstructing the Cognitive World*. Cambridge, MA: MIT Press.